Improving PIE's performance over high-delay paths

Nicolas Kuhn[†,⋆] and David Ros[‡]
[†]Institut Mines-Télécom / Télécom Bretagne, IRISA, 2 rue de la Châtaigneraie, 35576 Cesson Sévigné, France
[⋆]Centre National d'Etudes Spatiales (CNES), 18 Avenue Edouard Belin, 31400 Toulouse, France
[‡]Simula Research Laboratory, P.O. Box 134, 1325 Lysaker, Norway

arXiv:1602.00569v1 [cs.NI] 1 Feb 2016

***Abstract*—Bufferbloat is excessive latency due to over-provisioned network buffers. PIE and CoDel are two recently proposed Active Queue Management (AQM) algorithms, designed to tackle bufferbloat by lowering the queuing delay without degrading the bottleneck utilization. PIE uses a proportional integral controller to maintain the average queuing delay at a desired level; however, large Round Trip Times (RTT) result in large spikes in queuing delays, which induce high dropping probability and low utilization. To deal with this problem, we propose Maximum and Average queuing Delay with PIE (MADPIE). Loosely based on the drop policy used by CoDel to keep queuing delay bounded, MADPIE is a simple extension to PIE that adds deterministic packet drops at controlled intervals. By means of simulations, we observe that our proposed change does not affect PIE's performance when RTT < 100 ms. The deterministic drops are more dominant when the RTT increases, which results in lower maximum queuing delays and better performance for VoIP traffic and small file downloads, with no major impact on bulk transfers.**

***Keywords*—Active queue management; Internet latency; PIE.**

I. Introduction

ACTIVE queue management (AQM) schemes can be introduced in network routers with the goal of controlling the amount of buffering and reducing the loss synchronization. Large buffers and the absence of AQM deployment have resulted in huge latencies; this problem is known as bufferbloat [1]. Since AQM can control the queuing delay, its deployment "can significantly reduce the latency across an Internet path" [2]. The first AQM proposals, such as Random Early Detection (RED) [3], dating back more than a decade, have been reported to be usually turned off, mainly because of the difficulty to tune their parameters. Even if Adaptive RED (ARED) [4] was proposed to ease the parameterization of RED, it was designed to control the buffering when traffic is composed mainly of TCP flows. Proportional Integral controller Enhanced (PIE) [5] and Controlled Delay (CoDel) [6] are two recent AQM schemes that have been designed to tackle bufferbloat by lowering the queuing delay while addressing RED's stability issues and considering the presence of transports that do not react to congestion signals, such as UDP.

PIE and CoDel share two main concepts that can be mapped into two algorithm parameters: (1) a *target delay* (τ) represents the acceptable standing queuing delay above which an AQM drops packets more aggressively; (2) an update *interval* (λ) represents the reactivity of an AQM. These two parameters have different usages in the two algorithms. In CoDel, τ embodies an upper bound on the allowed queuing delay; if the minimum queuing delay over an interval of duration λ is higher than τ, then a packet is dropped with probability 1, else no packet is dropped. PIE uses τ to increase or decrease a dropping *probability*, based on the deviations of estimated queuing delay from such target delay: τ is therefore the desired average queuing delay.

CoDel has been shown to have auto-tuning issues and its performance is sensitive to the traffic load [7]. Also, its default 5 ms of maximum allowed queuing delay can be damaging for low-speed bottlenecks [8] and its interval value is based on the assumption that the Round Trip Time (RTT) is 100 ms [9]. On the other hand, PIE has been shown to be less sensitive to traffic loads [7], its default target delay of 20 ms should be less problematic with low capacity bottlenecks, and it does not make assumptions on the RTT. However, in this paper, we show that PIE is sensitive to the RTT, as we observe wide oscillations in queuing delay when the RTT increases. This results in temporarily high maximum queuing delay, high dropping probability, and low bottleneck utilization.

To reduce the RTT sensitivity of PIE and improve the performance of latency sensitive applications over large RTT paths (e.g., rural broadband or satellite access), our proposal, Maximum and Average queuing Delay with PIE (MADPIE) extends PIE by adding deterministic drops to prevent the queuing delay from growing beyond a critical value, loosely mimicking CoDel's drop policy.

The rest of this article is organized as follows. Section II details the MADPIE algorithm. In Section III, by means of simulations we illustrate the issues that PIE faces when the RTT increases, and how the deterministic drops in MADPIE help to correct those issues. Section IV provides an evaluation of the trade-off between allowing more bandwidth for bulk transfers and improving the performance of latency sensitive applications with MADPIE and PIE, as opposed to DropTail. Section V compares CoDel, PIE and MADPIE when the flows sharing the bottleneck do not have the same RTT. Finally, Section VI concludes this work.

II. Adding deterministic drops to PIE

PIE drops an incoming packet when $p \leq p_{drop}$, where p is drawn at random from a uniform distribution in $[0, 1]$, and p_{drop} is an internal variable updated every $\lambda = 30$ ms according to:

$$p_{drop} \leftarrow p_{drop} + \alpha \times (E[T] - \tau) + \beta \times (E[T] - E[T]_{old}). \quad (1)$$

$E[T]$ and $E[T]_{old}$ represent the current and previous estimation of the queuing delay. τ is PIE's target delay. α determines how the deviation of current queuing delay from τ affects the drop probability, whereas β exerts additional adjustments depending on whether the queuing delay is trending up or down.

MADPIE uses the same random drop policy as PIE, the only difference between the two algorithms being that we add a deterministic drop policy. MADPIE requires only one additional parameter: the queuing delay τ_{DD} above which deterministic drops occur. An indicator variable p_{max}, initialized to 0, tells whether a packet must be dropped ($p_{max} = 1$) or not ($p_{max} = 0$) by the deterministic policy. Every λ, if the estimated queuing delay is $> \tau_{DD}$, p_{max} is set to 1. Then, if a packet is *not* dropped nor marked by the random drop algorithm and $p_{max} = 1$, then a packet is dropped or marked and p_{max} is reset to 0. Thus, there can be a maximum of one deterministic drop every λ.

III. Proof of Concept

The aim of this section is to illustrate how MADPIE's behaviour differs from that of PIE when the RTT increases.

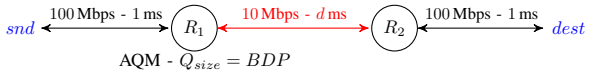

Fig. 1. Topology used to prove the MADPIE concept.

Fig. 1 presents the topology used in this section. The one-way delay of the bottleneck link is set to $d = 48$ ms or $d = 248$ ms (which corresponds to a base RTT[1] of 100 ms and 500 ms, respectively). The queue size at R_1 is set to the Bandwidth-Delay Product (BDP). The AQM introduced at R_1 is either PIE ($\tau = 20$ ms and $\lambda = 30$ ms) or MADPIE ($\tau = 20$ ms, $\lambda = 30$ ms, $\tau_{DD} = 30$ ms). We simulate 10 TCP bulk flows from snd to $dest$, using CUBIC as congestion control policy, for 300 s. The Initial congestion Window (IW) is set to 10 packets and the SACK option is enabled. The flows randomly start in $[0; 1]$ s. All TCP variants used in this article were provided by the NS-2 TCP Linux module updated to linux kernel version `3.17.4`.[2].

In Fig. 2, we present the Cumulative Distribution Function (CDF) of the queuing delay (measured per packet) and the bottleneck utilization (sampled every second). When the RTT is 100 ms, apart from the maximum queuing delay that is slightly lower with MADPIE than with PIE (≈ 100 ms with PIE, ≈ 80 ms with MADPIE), there is no noticeable performance difference between MADPIE and PIE. When the RTT is 500 ms, Fig. 2a and 2c show that for 20 % of the samples, the queuing delay is higher than 20 ms with MADPIE whereas it is higher than 30 ms with PIE. Also, with MADPIE as opposed to with PIE, the maximum queuing delay is reduced by ≈ 60 ms. It is worth pointing out that 90 % of the samples show a queuing delay lower than 30 ms (that is τ_{DD}) with MADPIE as opposed to 50 ms with PIE. Fig. 2b and 2d show

[1]That is, the minimum RTT, without any queuing delays.

[2]More details at: http://heim.ifi.uio.no/michawe/research/tools/ns/index.html

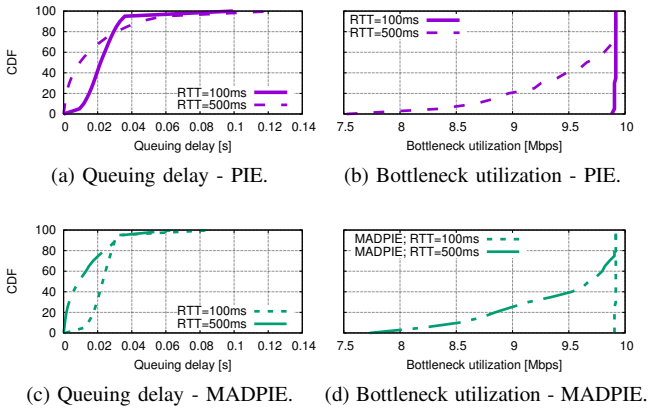

(a) Queuing delay - PIE. (b) Bottleneck utilization - PIE.

(c) Queuing delay - MADPIE. (d) Bottleneck utilization - MADPIE.

Fig. 2. Bottleneck utilization and queuing delay.

that this latency reduction does not induce a lower bottleneck utilization.

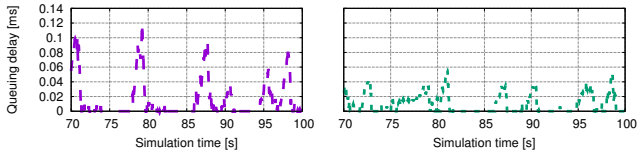

(a) PIE (b) MADPIE

Fig. 3. Evolution of queuing delay over time with $RTT = 500$ ms.

This latency reduction provided by MADPIE can be further explained by looking at the queuing delay evolution in Fig. 3. With PIE, a higher RTT results in wider oscillations in queuing delay: as the queuing delay gets much higher than τ, the dropping probability increases in order to maintain a lower queuing delay. This however results in a momentarily empty buffer. PIE's burst allowance of 100 ms lets the queuing delay to frequently grow above 100 ms, as the buffer was previously empty. With MADPIE, it is possible to initially allow the same bursts, but the deterministic drops would then prevent an excessive growth of both the queuing delay and the drop probability if the buffer is frequently empty then full, which is what happens when the RTT is 500 ms.

To better understand how MADPIE's behaviour differs from that of PIE when the RTT increases, we look at the contribution of random and deterministic drops to the overall drop rate. Let us denote by n_{DD}, n_{RD}, n_{BO} and n_{tot} the number of drop events induced by a deterministic drop (only MADPIE), a random drop (PIE and MADPIE), a buffer overflow (PIE and MADPIE) and the total number of drops, respectively. Let $r_x = n_x/n_{tot}$ be the proportion of drop events of type x. Fig. 4 shows r_{RD} and r_{DD} as a function of the queuing delay when the drop occured. As one example (dashed lines in Fig. 4), when the RTT is 500 ms and the AQM is MADPIE, $r_{DD} \approx 70$ % and when the deterministic drops occured, the

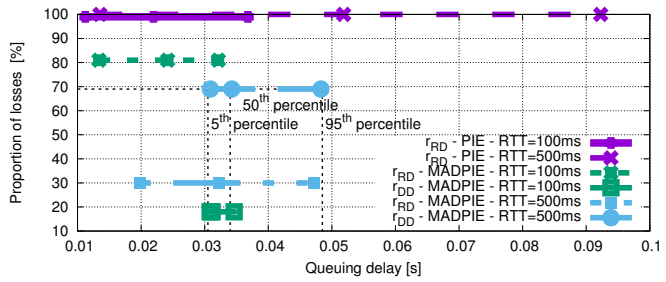

Fig. 4. Queuing delay and proportion of losses.

5 % percentile of the queuing delay was ≈ 30 ms, the 50 % percentile ≈ 34 ms and the 95 % percentile ≈ 48 ms. With PIE, most of the drops are induced by the AQM algorithm and not by buffer overflow and, when the RTT is 500 ms, the queuing delay raises up to more than 90 ms. With MADPIE, when the RTT is 100 ms, the random drop part of MADPIE is responsible for more than 80 % of the drops, whereas when the RTT is 500 ms, the deterministic part of MADPIE is responsible for around 70 % of the drops with a consequent queuing delay reduction.

IV. Performance of MADPIE with a traffic mix

We compare now the performance of DropTail (DT), PIE and MADPIE when the traffic comes from a mix of various applications.

A. Traffic and topology

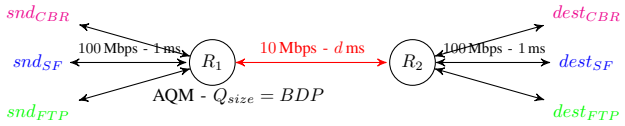

Fig. 5. Topology and traffic mix used to evaluate MADPIE.

Fig. 5 presents the topology used in this section. The one-way delay d of the bottleneck link is set to 48 ms, 148 ms or 248 ms (*i.e.* a base RTT of 100 ms, 300 ms or 500 ms). The queue size at R_1 is set to the BDP. The AQM introduced at R_1 is either PIE ($\tau = 20$ ms and $\lambda = 30$ ms) or MADPIE ($\tau = 20$ ms, $\lambda = 30$ ms, $\tau_{DD} = 25$ ms).

Between snd_{CBR} and $dest_{CBR}$, there are N_{CBR} Constant Bit-Rate (CBR) UDP flows with a sending rate of 87 kbps and a packet size of 218 B. The intent is to model Voice-over-IP (VoIP) or gaming traffic, such as in [10, p. 17]. Between snd_{SF} and $dest_{SF}$, N_{SF} flows transfer files of S kB ($S \in \{15; 44; 73; 102\}$). When a download is finished, a new random value is taken for S and another download starts after τ seconds, with τ randomly generated according to an exponential law of mean 9.5 s. This traffic lets us assess the benefits of using MADPIE for short flows. Between snd_{FTP} and $dest_{FTP}$, N_{FTP} TCP bulk flows are generated. TCP flows use CUBIC congestion control, and TCP options are the same as those specified in § III. All the flows randomly start between 0 and 1 s. Each run lasts 100 s and is repeated 20 times with independent seeds. The metrics are sampled every second (except for the queuing delay and the one way delay that are sampled per-packet). We choose to present the results with $N_{CBR} = 4$, $N_{SF} = 20$ and $N_{FTP} = 10$, as this traffic mix stresses both PIE and MADPIE.

B. CBR traffic

The performance for CBR traffic is shown in Fig. 6. Fig. 6a explains how to interpret Fig. 6b, 6c and 6d. We present the average cumulative goodput as a function of the queuing delay, as advised in [11].

The results with DT, shown in Fig. 6b, illustrate that queuing delay can be very high, impacting latency-sensitive applications (the higher percentiles for the queuing delay when the RTT is 500 ms are not shown as they do not fit in the current scale). The goodput may sometimes be over 87 kps as delayed packets at the bottleneck queue may arrive in bursts at the receiver. The comparison of the results of PIE (in Fig. 6c) and MADPIE (in Fig. 6d) confirms that when the RTT is 100 ms, MADPIE does not differ much from PIE. When the RTT increases, the deterministic drops induced by MADPIE allow a reduction in the experienced queuing delay of 5 ms for the 75^{th} percentile and of 30 ms for the 95^{th} percentile, without noticeable impact on the goodput.

C. Small-file downloads

We represent in Fig. 7 the download time of files of various sizes, with and without AQM schemes; the boxplots show the 5^{th}, 25^{th}, 75^{th} and 95^{th} percentiles; the line in the middle of the box is the median. With DT, the download time is higher than with any of the two AQMs for every file size and RTT considered. Comparison of the results with MADPIE and PIE shows that MADPIE reduces worst-case transmission times. For example, with MADPIE as compared with PIE, (1) the 95^{th} percentile of the download time for 73 Kb is reduced by ≈ 700 ms when $RTT = 300$ ms; (2) the 75^{th} percentile of the download time for 102 Kb is reduced by ≈ 500 ms when $RTT = 500$ ms. This can be explained by the fact that, with MADPIE, the few packets that compose a short file transfer have a lower probability of experiencing high queuing delays, and of arriving at the queue when the random-drop probability is high (hence suffering losses in a burst).

D. Bulk flows

Fig. 8 shows the CDF of the goodput for the bulk flows. With DT, the impact of the RTT can hardly be noticed. Without AQM at R_1, the bottleneck utilization is higher than with any of PIE or MADPIE. With the latter, when the RTT is large the gain for latency sensitive applications comes at the expense of a small degradation in goodput for bulk flows.

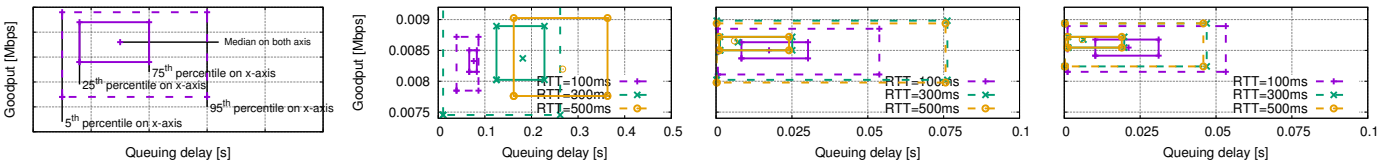

(a) Example. (b) DT. (c) PIE. (d) MADPIE.

Fig. 6. Average goodput and queuing delay for the CBR traffic.

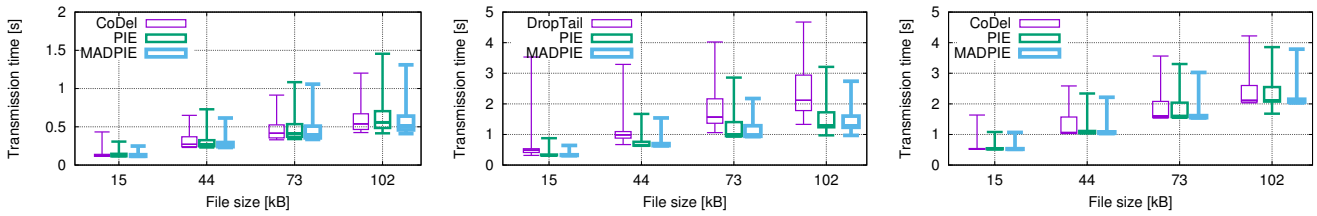

(a) $RTT = 100$ ms. (b) $RTT = 300$ ms. (c) $RTT = 500$ ms.

Fig. 7. Small file download time.

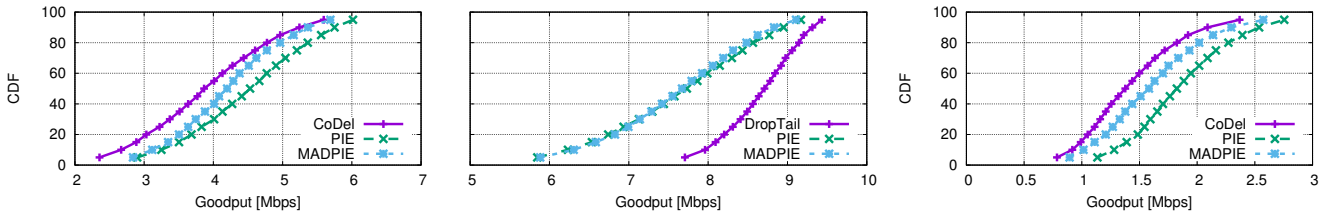

(a) $RTT = 100$ ms. (b) $RTT = 300$ ms. (c) $RTT = 500$ ms.

Fig. 8. Goodput of the bulk flows.

V. Traffic mix and RTT mix

Previous section focused on highlighting the difference between PIE and MADPIE; for the sake of clarity, we preferred showing results only with PIE and MADPIE. Since this section consists more of a performance analysis, rather than of a behaviour analysis, we also show results with CoDel, to compare the benefits given by introducing any of the considered AQM. We consider flows that do not face the same RTT to assess the benefits of using a scheme that is less sensitive to RTT. We do not present the results with DropTail, since § IV showed that it results in poor performance for latency sensitive applications when the RTT increases.

A. Traffic and topology

Figure 9 presents the topology used in this section. The traffic generated, the parameterization of PIE and MADPIE, and the methodology (number of runs, duration of each run) have the same characteristics as in § IV, the only difference being the number of flows that is the following: $N_{CBR-100} = N_{CBR-500} = 4$, $N_{SF-100} = N_{SF-500} = 20$ and $N_{FTP-500} = N_{FTP-500} = 2$, where N_{X-100} represents

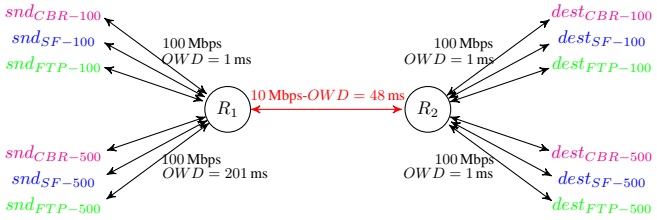

Fig. 9. Topology and traffic mix used to evaluate the RTT sensitivity

the number of flows for the application of type X on the path with 100 ms of RTT. As the RTT of the paths are not the same, the queue size at R_1 is set to the BDP of the higher RTT.

B. CBR traffic - between snd_{SF-X} and $dest_{SF-X}$

The results for the CBR traffic are shown in Figure 10 and we use the same representation as in Figure 6a. Due to the presence of flows that experience an RTT of 500 ms, the flows with an RTT of 100 ms face a queuing delay that momentarily rises above 60 ms with PIE (*i.e.* 55 % of the

one way delay). Since the default *target delay* of CoDel is 5 ms, the allowed queuing delay is lower than with PIE and MADPIE, for which the *target delay* is set to 20 ms. By default, CoDel would maintain a lower queuing delay than PIE. Also, because CoDel uses deterministic drops, the queuing delay can not rise much higher before the first drops are applied. These results show that MADPIE takes the best of the two schemes: with MADPIE, the median queuing delay is lower than with PIE and the queuing delay is kept under control, and the higher percentiles of the queuing delay are close to those with CoDel. While with PIE, the queuing delay momentarily increases above 60 ms (*i.e.* 55 % of the RTT), with MADPIE, the introduction of the flows that experience an RTT of 500 ms has less impact: the queuing delay momentarily rises above 40 ms, that is 45 % of the RTT. MADPIE provides a latency reduction of $\approx$ 13 % for the 75^{th} percentiles and of $\approx$ 21 % for the 95^{th} percentiles and the performance is close to that of CoDel.

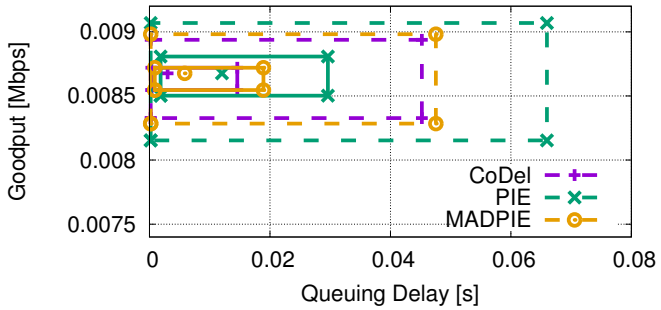

Fig. 10. Goodput and queuing delay - CBR traffic for which $RTT = 100$ ms

C. Small files download - between snd_{SF-X} and $dest_{SF-X}$

We show in Figure 11 the download time of small files for the flows that experience an RTT of 100 ms (Figure 11a) and of 500 ms (Figure 11b).

For small file downloads over the path with a base RTT of 100 ms, the download time is lower than in § IV-C as there is much less concurrent traffic: in this section, there are 4 bulk flows as opposed to 10 in § IV. The 5^{th}, 25^{th}, 50^{th} and 75^{th} percentiles are quite the same whether the AQM is PIE, MADPIE or CoDel; the 75^{th} percentile is slightly lower with MADPIE. The 95^{th} percentile is always lower with MADPIE. When the file size is larger than 73 kB, the 95^{th} percentile is slightly lower with CoDel or MADPIE than with PIE, probably thanks to the fact that if the AQM drops a packet, the queuing delay experienced by the retransmission would be lower. For small file downloads that experience a base RTT of 500 ms, the performance is quite the same with CoDel and PIE: with PIE, the bottleneck utilisation is higher (see § V-D for more details) and with CoDel the queuing delay is lower (see § V-B for more details). MADPIE takes the best of both schemes and thus, provides lower download times for small files.

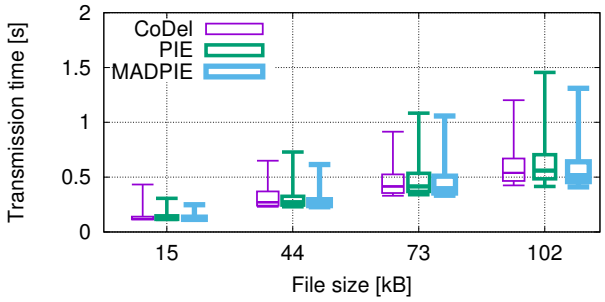

(a) $RTT = 100$ ms

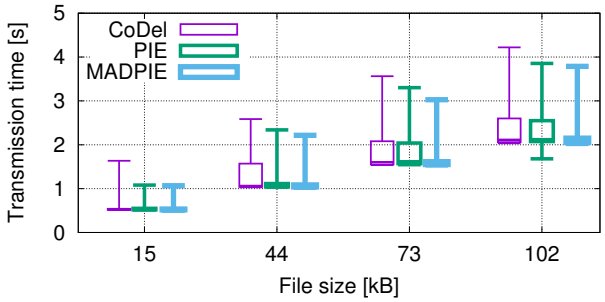

(b) $RTT = 500$ ms

Fig. 11. Small file download time

D. Bulk flows - between snd_{FTP-X} and $dest_{FTP-X}$

We show in Figure 12 the goodput of bulk transfers for flows that experience an RTT of 100 ms (Figure 12a) and of 500 ms (Figure 12b). CoDel shows a lower goodput than PIE, which is due to its lower *target delay*: PIE allows more buffering. As seen in § IV-D, MADPIE slightly reduces the bottleneck utilization for the bulk flows that experience an RTT of 100 ms. The same happens for the flows that experience an RTT of 500 ms. With MADPIE, the resulting goodput is a trade-off between CoDel and PIE.

E. Discussion

The benefits of using MADPIE instead of CoDel may not be clear, however we advice to deploy MADPIE instead of CoDel. Indeed, in [7], CoDel has been shown to have auto-tuning issues and its performance are sensitive to the load of traffic. Also, 5 ms of maximum allowed queuing delay can be damaging for bottlenecks of 2 Mbps [8] and its interval value is based on the assumption that the RTT is 100 ms [9], which is not the case for rural broadband networks. On the contrary, our algorithm clearly improves the performance of PIE when the RTT is higher than 300 ms for various types of traffic and does not affect the performance of PIE for lower RTTs. The deployment issues of CoDel mentioned earlier in this section are solved with PIE as: (1) it is less sensitive to traffic loads [7]; (2) with 20 ms of targeted queuing delay, we expect less issues with low capacity bottlenecks; (3) it does not make assumptions on the RTT.

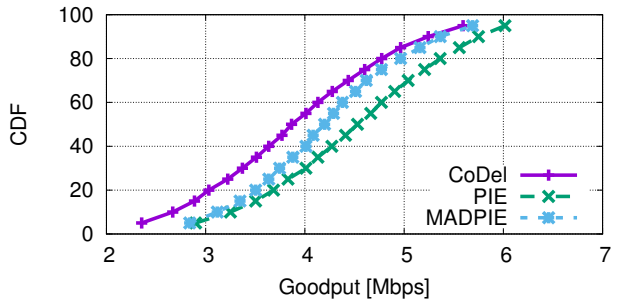

(a) $RTT = 100\,\text{ms}$

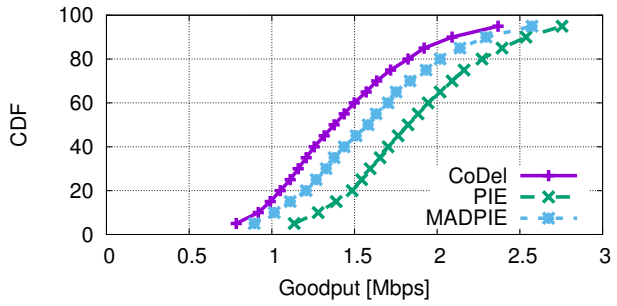

(b) $RTT = 500\,\text{ms}$

Fig. 12. Goodput of the bulk flows

VI. Conclusion

In this paper, we have proposed MADPIE, a simple change to the PIE algorithm that makes it less dependent on path RTTs in lightly-multiplexed scenarios. MADPIE extends PIE by adding, on top of the random drops, a deterministic drop policy loosely based on CoDel's. The proportion of deterministic drops increases when the RTT increases. MADPIE can both keep the same target delay as PIE and reduce the maximum queuing delay, making the goodput of bulk flows close to the one achieved with PIE, guaranteeing lower queuing delay for VoIP-like traffic and reducing the download time of small files.

We do not claim that our proposal is the only, or best, way of tuning or adapting PIE. However, it is a very simple addition to PIE's code (a handful of lines, in our ns-2 implementation) that can complement specific parameter tunings, and its impact on the performance of PIE seems negligible when RTTs are not large (i.e, outside the operating conditions for which it has been conceived).

Acknowledgement

This work was partially funded by the European Community under its Seventh Framework Programme through the Reducing Internet Transport Latency (RITE) project (ICT-317700).

References

[1] J. Gettys, "Bufferbloat: Dark Buffers in the Internet," *IEEE Internet Computing*, vol. 15, no. 3, pp. 96–96, May 2011.

[2] F. Baker and G. Fairhurst, "IETF Recommendations Regarding Active Queue Management," Best Current Practice RFC 7567, IETF, Jul. 2015. [Online]. Available: http://tools.ietf.org/html/rfc7567

[3] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Transactions on Networking*, vol. 1, no. 4, pp. 397–413, Aug. 1993.

[4] S. Floyd, R. Gummadi, and S. Shenker, "Adaptive RED: An algorithm for increasing the robustness of RED's active queue management," ICIR, Technical report, Aug. 2001. [Online]. Available: http://www.icir.org/floyd/red.html

[5] R. Pan, P. Natarajan, C. Piglione, M. Prabhu, V. Subramanian, F. Baker, and B. VerSteeg, "PIE: A lightweight control scheme to address the Bufferbloat problem," in *Proceedings of IEEE HPSR*, Taipei, Jul. 2013.

[6] K. Nichols and V. Jacobson, "Controlling Queue Delay," *ACM Queue*, vol. 10, no. 5, pp. 20:20–20:34, May 2012.

[7] I. Järvinen and M. Kojo, "Evaluating CoDel, PIE, and HRED AQM techniques with load transients," in *Proceedings of IEEE LCN*, Edmonton, Sep. 2014, pp. 159–167.

[8] C. Kulatunga, N. Kuhn, G. Fairhurst, and D. Ros, "Tackling Bufferbloat in Capacity-limited Networks," in *24th European Conference on Networks and Communications (EuCNC)*, Paris, Jun. 2015.

[9] K. Nichols, V. Jacobson, A. McGregor, and J. Iyengar, "Controlled delay active queue management," Internet Draft draft-ietf-aqm-codel, work in progress, Apr. 2015. [Online]. Available: https://tools.ietf.org/html/draft-ietf-aqm-codel

[10] G. White and D. Rice, "Active Queue Management Algorithms for DOCSIS 3.0," Cable Television Laboratories, Tech. Rep., Apr. 2013. [Online]. Available: http://www.cablelabs.com/wp-content/uploads/2013/11/Active_Queue_Management_Algorithms_DOCSIS_3_0.pdf

[11] N. Kuhn, P. Natarajan, N. Khademi, and D. Ros, "AQM characterization guidelines," Internet Draft draft-ietf-aqm-eval-guidelines, work in progress, Jul. 2015. [Online]. Available: https://tools.ietf.org/html/draft-ietf-aqm-eval-guidelines